Lifting Micro-Update Models from RTL for Formal Security Analysis

Adwait Godbole
adwait@berkeley.edu
University of California, Berkeley
Berkeley, CA, USA

Kevin Cheang*
kcheang@berkeley.edu
University of California, Berkeley
Berkeley, CA, USA

Yatin A. Manerkar
manerkar@umich.edu
University of Michigan
Ann Arbor, MI, USA

Sanjit A. Seshia
sseshia@berkeley.edu
University of California, Berkeley
Berkeley, CA, USA

Abstract

Hardware execution attacks exploit subtle microarchitectural interactions to leak secret data. While checking programs for the existence of such attacks is essential, verification of software against the full hardware implementation does not scale. Verification using abstract formal models of the hardware can help provide strong security guarantees while leveraging abstraction to achieve scalability. However, hand-writing accurate abstract models is tedious and error-prone. Hence, we need techniques to generate models that enable sound yet scalable security analysis automatically.

In this work, we propose *micro-update models* as a modelling framework that enables sound and abstract modelling of microarchitectural features. We also develop algorithms to generate micro-update models from RTL semi-automatically. We implement our modelling and generation framework in a prototype tool called PAUL. We evaluate our approach by synthesizing micro-update models for the `Sodor5Stage` processor and components from the `cva6` (Ariane) processor. We demonstrate how these models can be generated hierarchically, thus increasing scalability to larger designs. We observe up to 8× improvement in run time when performing analysis with the generated models as compared to the source RTL.

CCS Concepts • **Security and privacy** → **Hardware reverse engineering**; **Side-channel analysis and countermeasures**; *Formal security models*; *Logic and verification*; • **Hardware** → *Functional verification*; • **Theory of computation** → **Verification by model checking**; *Models of computation*; • **Software and its engineering** → *Formal methods*;

ACM Reference format:
Adwait Godbole, Kevin Cheang, Yatin A. Manerkar, and Sanjit A. Seshia. 2024. Lifting Micro-Update Models from RTL for Formal Security Analysis. In *Proceedings of 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2, La Jolla, CA, USA, April 27-May 1, 2024 (ASPLOS '24)*, 14 pages.
https://doi.org/10.1145/3620665.3640418

1 Introduction

The landscape of hardware execution attacks has evolved since Spectre [44] and Meltdown [46]. Recent attacks exploit deeper microarchitectural state such as store buffers and line-fill-buffers [16, 62, 72, 73]. Checking software for the presence of these attacks requires analyses that take this microarchitectural state into account. However, owing to microarchitectural complexity, verification of software programs directly against the source RTL does not scale with the program being verified. Abstract formal models that only represent parts of the microarchitecture relevant to the attack while eliding the rest are simpler and more interpretable than the RTL. Verifying SW against such models instead of the source HW can improve scalability. However, accurate modelling is necessary for preserving strong guarantees; imprecise models can lead to software being deemed secure when in reality it is not.

Several recent works apply formal methods to verify software against hardware attacks [18, 29, 30, 55], and to identify new ones [23, 70]. These works perform verification against *manually crafted* abstract models that preserve only the vulnerability-relevant components of the HW design. However, manually writing models while ensuring the accuracy of preserved components is tedious and error-prone, and requires a deep understanding of the microarchitecture. While a security analyst can indicate *which* source design components they want to perform software-side security analysis against, it is hard to manually specify *how* these components should behave so that the model is sufficiently accurate. Thus, it is desirable to have a technique that, given some *signals-of-interest*, generates a formal model that accurately captures the behaviours of these signals w.r.t. the source RTL.

In this work, we propose ***micro-update models*** as a modelling formalism that can be used in security analysis of software running on the source RTL design. To address the challenges involved in hand-writing models, we develop a framework to mostly automatically ***generate micro-update models*** from RTL designs. The generated model achieves abstraction by capturing signals-of-interest specified by the user, eliding the rest, and achieves accuracy by ensuring functional equivalence of these signals w.r.t. the source RTL. This guarantees soundness, i.e. security properties over the signals-of-interest (e.g. information-flow/non-interference (NI) [21, 27]) that hold on the lifted model also holds on the RTL.

Imperative modelling of the functional behaviour of the signals-of-interest in the micro-update model: (a) aids correct-by-construction generation, since the model can be checked against the source RTL (e.g. with a functional equivalence proof) and (b) enables sound verification of software w.r.t. security properties (e.g. NI). In contrast, ISA-based models (e.g. [19, 28]), are equivalent only w.r.t. the *architectural state*, and lack the microarchitectural detail necessary for security analysis. On the other hand, approaches such as μspec [48, 49] formulate *axiomatic* ordering models that abstract away functional behaviours. Due to the lack of imperatively modelled state, μspec requires global invariants for checking accuracy with

*Work was performed when Kevin was at UC Berkeley.

the source RTL. Defining these invariants is challenging [35, 51]. Secondly, properties such as NI [21, 27], which requires a model that preserves functional behaviours cannot be expressed with these models. In conclusion, imperatively modelling functional behaviours of the signals-of-interest makes checking alignment with source RTL, and security properties on SW easier with the micro-update model.

Our micro-update model generation framework is based on *lifting*, which aims to extract high-level code from low-level source implementations while maintaining important properties with respect to the source. While lifting has been demonstrated in the PL and systems domains [3, 4, 41, 65], we extend it to hardware. We demonstrate how lifting can be performed *hierarchically* based on *oracle-guided inductive synthesis* [40, 63], thereby increasing its scalability. Our techniques aid in generating micro-update models from RTL, which then can be used for software security analysis.

Contributions. Our contributions are as follows:

- **Micro-update models as a new formalism:** We propose micro-update models as a formalism for abstract modelling of hardware designs. They provide finer temporal resolution than ISA-level models and preserve functional behaviour. Thus, they provide accurate abstractions of the hardware against which software can be verified for microarchitectural attacks.
- **Synthesis techniques for micro-update models:** We develop a methodology to generate micro-update models from RTL using a novel instantiation of *oracle-guided synthesis*[1] [39, 40, 60, 63]. We show how one can generate models *hierarchically*, improving the scalability of synthesis, and provably maintaining a notion of equivalence with the RTL.
- **Empirical Demonstration:** We evaluate our approach on the `Sodor5Stage` processor [71] and components from the `cva6` processor [1] by generating micro-update models for them. We demonstrate how applying the generated models to perform security analysis of software can lead to performance improvements over the source RTL.

Outline. In §2 we motivate and contrast the features of our modelling approach with existing ones through an example. In §3 we describe the structure of our models and their semantics. In §4 we describe the techniques that we use to generate our models. We discuss the experimental evaluation in §5. Finally, §6 discusses limitations, §7 is related work and §8 concludes.

2 Motivating Example

2.1 Example: modified `Sodor5Stage` core

`Sodor5Stage` [71] is a 5-stage in-order processor implementing the RV32UI instruction set [75]. We augment this processor design with a new component, called a load buffer (LB) as shown in Fig. 1. The LB caches the most recently loaded value in its `data` field. It also maintains an address field and a `valid` bit marking whether the entry in the buffer is not outdated. This cached value can be consumed by a subsequent load if its address matches and the entry is valid: *canLoad* ≡ (`loadAddr == addr`) ∧ `valid`.

The LB entry is flushed by any subsequent store instruction to prevent reads from outdated entries. This feature does not change the ISA-level behaviour of the processor (program counter, register file and memory); its effects are purely microarchitectural. Now we discuss how a hypothetical attack can leak information by using the LB as a side channel.

[1]Oracles are humans/tools that provide guidance to the synthesis engine in a specified format through a query-response interface.

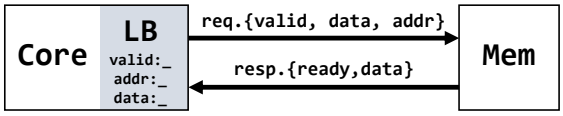

Figure 1. Schematic for the `Sodor5Stage` processor. Our modification, the load buffer (LB), is shown in grey.

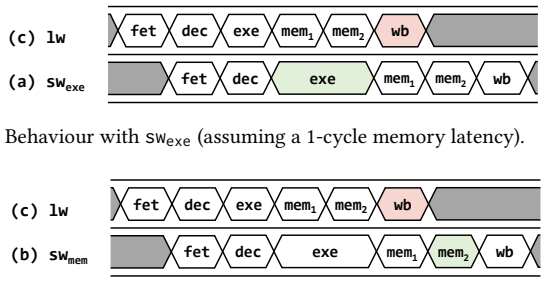

Figure 2. Different microarchitectural behaviours (with a 1-cycle memory latency) for different implementation choices of `sw`.

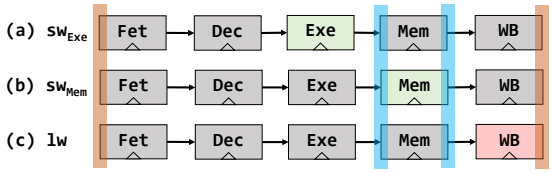

Figure 3. Pipeline stages for two possible implementations for the `sw` instruction: in (a) the LB flush happens during the `exe` stage, while in (b) it happens during the `mem` stage; (c) illustrates that the refill of the LB due to the `lw` instruction takes place in its `wb` stage.

2.2 The `lb` optimization vulnerability

Microarchitectural features such as the LB, can have several possible implementations. In Fig. 3(a,b) we consider two implementations for `sw` which differ in the exact cycle at which the LB is invalidated by the store. In Fig. 3(a) the invalidation happens during the `exe` stage of the `sw` instruction, while in (b) it does so during the `mem` stage. We call these implementations $\texttt{sw}_{\texttt{exe}}$ and $\texttt{sw}_{\texttt{mem}}$ respectively. A `lw` instruction loads from memory in the `mem` stage and refills the LB in the `wb` stage in both cases (Fig. 3(c)).

While architecturally invisible, such implementation differences crucially affect security analysis. To see this let us consider the effect of executing a load consecutively followed by a store, and compare them across the $\texttt{sw}_{\texttt{exe}}$ and $\texttt{sw}_{\texttt{mem}}$ implementations. In the case of `lw`; $\texttt{sw}_{\texttt{exe}}$ (Fig. 2 top), the LB has a valid entry at the end of execution (refilled due to the `lw` instruction). On the other hand, in the case of `lw`; $\texttt{sw}_{\texttt{mem}}$ (Fig. 2 bottom), the LB entry is flushed by the `sw` *after* the refill caused by `lw`. Consequentially, `lw`; $\texttt{sw}_{\texttt{exe}}$ leads to an LB which is tainted (due to the `lw`), while `lw`; $\texttt{sw}_{\texttt{mem}}$ does not.

A timing-based side-channel (e.g. Prime+Probe [25, 47]), can allow an attacker to infer the contents of the cache-like LB through a timing measurement. In such a case, the `lw`; $\texttt{sw}_{\texttt{exe}}$ can lead to secret data (e.g. private keys) getting revealed to the attacker. On the other hand, since the `lw`; $\texttt{sw}_{\texttt{mem}}$ sequence leads to a *sanitized* LB, (i.e. the LB is untainted), a timing measurement cannot reveal victim secrets. Such microarchitectural details can render software insecure under certain implementations, while secure under others. Hence analyses checking software security must operate on models that expose microarchitectural detail.

2.3 Instruction-level approaches

Models that capture ISA-level behaviours [14, 19, 28, 66, 74] are precise only with respect to software-visible architectural state (e.g. program counter, register file, memory, and CSRs). This is also true of instruction-level-abstraction (ILA) based approaches [36, 38, 81] which model accelerators. However, hardware attacks manifesting at the microarchitectural (e.g., [16, 44, 46, 62, 72, 73]) level exploit software-invisible features (e.g. caches, predictors). Consequentially, ISA-level models are restricted to certain software-level abstractions of security properties for these attacks. One example is the constant time software contract, which requires instruction counts and memory access sequences to be independent of any secret information. While constant-time can be expressed over the architectural state, it is sensitive to the microarchitectural implementation [17]. Abstracting away microarchitectural detail can also lead constant time to impose overly strong requirements on software. In the LB example (Fig. 3), the instruction sequence `lw`; $\texttt{sw}_{\texttt{mem}}$ would be *secure* (assuming the `lb_addr` is the observable state) even if the `lw` loaded data from a secret-dependent address. This is because the `sw` would flush the `lb_addr`. However, since the load address is secret dependent, this code is not constant-time.

ISA-level approaches also subscribe to a temporal granularity that corresponds to the isolated execution of individual software (assembly) instructions. To illustrate this, let $[\![i_1; i_2]\!]$ denote the effect (semantics) of the sequence of two instructions on the HW design. If one is only concerned with the architectural state, σ_{arch}, the effect of $[\![i_1; i_2]\!]$ is equivalent to the effect of these instructions executed in sequentially and *in isolation*, $[\![i_1]\!]$ followed by $[\![i_2]\!]$ (i.e., i_1 finishes execution before i_2 begins):

$$\sigma_{\text{arch}}\ [\![i_1; i_2]\!]\ \sigma'_{\text{arch}} \iff \sigma_{\text{arch}}\ [\![i_1]\!]\ \sigma''_{\text{arch}}\ [\![i_2]\!]\ \sigma'_{\text{arch}}$$

However, the effect of a stream of instructions with respect to the *microarchitectural* state $\sigma_{\mu\text{arch}}$ can be different than that of individual instructions operating in isolation:

$$\sigma_{\mu\text{arch}}\ [\![i_1; i_2]\!]\ \sigma'_{\mu\text{arch}} \not\iff \sigma_{\mu\text{arch}}\ [\![i_1]\!]\ \sigma''_{\mu\text{arch}}\ [\![i_2]\!]\ \sigma'_{\mu\text{arch}}$$

In the LB example, the semantics of both $\texttt{sw}_{\texttt{exe}}$ and $\texttt{sw}_{\texttt{mem}}$ are identical when executed in isolation. This is true *even with respect to the LB state*, as both flush the LB before finishing execution. However, the semantics of $\texttt{sw}_{\texttt{exe}}$; `lw` and $\texttt{sw}_{\texttt{mem}}$; `lw` are different owing to differences in microarchitectural implementation.

In summary, models that define semantics for isolated execution between instruction fetch and commit points (orange lines in Fig. 3) fall short when representing the microarchitecture. Thus, we need a model that accurately captures the microarchitectural interactions between instructions.

2.4 Axiomatic ordering-based approaches

Approaches such as μspec [48] model the microarchitecture as a set of *axiomatic ordering constraints* over the execution of a program. Executions satisfying these constraints are valid (can be observed) while others cannot. Since these constraints are over full executions, μspec can express behaviours of instructions executing simultaneously (as in a pipeline), and does not suffer from the isolation problem discussed earlier.

Ordering-based approaches often abstract away functional behaviour. However, functional behaviour is necessary to model low-level security-relevant microarchitecture, as we now discuss. Consider a victim program running on a processor with an LRU-cache. Such processors can leak victim secrets through the LRU state [79]. Detecting such vulnerabilities in the victim program requires security properties that precisely identify memory access sequences which lead to attacker-distinguishable LRU states. A non-interference (NI) [27, 69] property identifies such accesses by requiring that subsequent attacker-observable cache hit/miss outcomes should not depend on victim secrets. However, since expressing security properties such as NI requires a model that exposes the functional behaviour of signals, it is infeasible with ordering-only models. A micro-update model can represent the low-level LRU policy functionally, enabling NI-based analysis.

μspec-based security analysis approaches (e.g. [55, 70]) work around this by explicitly identifying execution fragments which are indicative of vulnerabilities, e.g., exploit patterns (§3A-II in [70]) or *transmitters* (§3.2.1 in [55]). While NI uniformly captures all possible vulnerable cases as a single property, these execution fragments (patterns/transmitters) must be manually enumerated. This requires a deep understanding of the microarchitecture and can result in false negatives if some vulnerable execution fragments are missed.

The second advantage of the micro-update model is that it is *operational*. That is, it represents executions as imperative transitions performed on some state at every step. This allows checking *per-step* equivalence (§4.2) with the RTL design using a model checker (e.g. SymbiYosys [20], JasperGold). Such checks are essential to ensure accuracy when lifting a model. Ordering-based models, on the other hand, are *axiomatic* and define behaviours in terms of declarative constraints. Checking these constraints against the RTL requires *global invariants* (i.e. invariants over full executions as opposed to individual steps). These invariants can be difficult to generate and check automatically (e.g. [35, 51]), which in turn makes checking accuracy with the source RTL more difficult.

In conclusion, the micro-update model (a) represents executions at a finer granularity than ISA models, (b) preserves functional behaviour enabling NI-based security analysis and (c) represents executions operationally, thus allows easier equivalence checking with the source RTL.

3 Micro-update Models

In this section, we illustrate the micro-update model through Figure 4. A micro-update model is defined over a set of signals, and has two kinds of elements: micro-updates and guards. A *micro-update* is an imperative code snippet that defines a functional operation performed on signals from the model. For example, the flush micro-update (`lb_flush`) in Fig. 4(B) operates on the LB (from Fig. 1): it clears the data and address fields and sets the valid bit to zero.

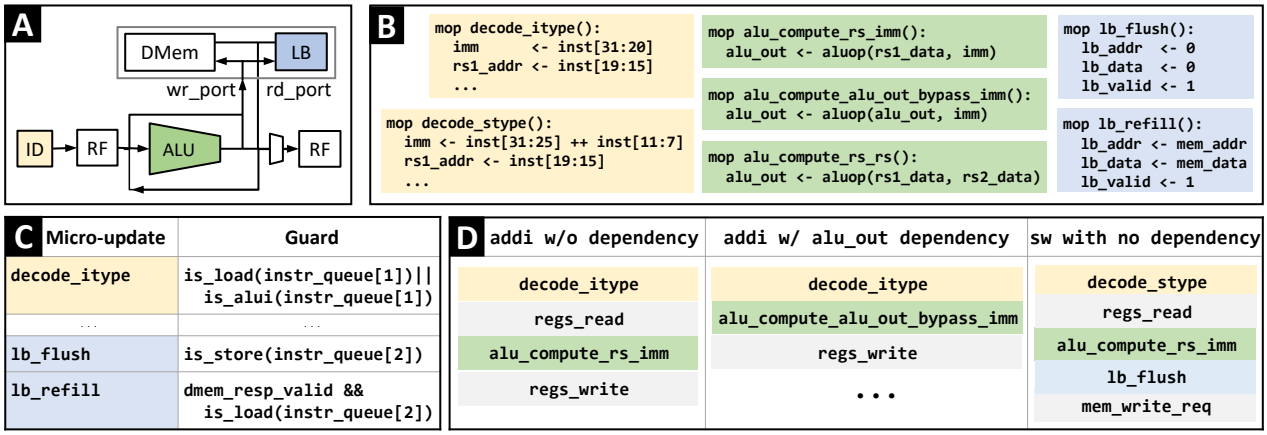

Figure 4. (A) A simplified processor with a LB and bypassing paths (greyed), **(B)** non-exhaustive set of micro-updates for this processor, **(C)** guards corresponding to these micro-updates, **(D)** sample instructions with some of the induced micro-updates.

Guards are Boolean predicates that control *when* a micro-update is triggered. Guards are defined over a set of *trigger signals*. For example in Fig. 4(C), the guard for `lb_flush` is true and `lb_flush` is executed at cycles where the instruction in the memory stage (`instr_queue[2]`) is a store. Trigger signals typically are top-level inputs to the source design. Signals from the model map to signals in the source RTL design. Thus, the model as a whole captures a slice of the design, as identified by these *signals-of-interest* (§3.1.1).

Example 3.1. Consider an `addi` instruction executing on the simple pipelined processor in Fig. 4(A). During execution `addi` interacts with several pipeline components (e.g. decode unit), by invoking multiple micro-updates. Each micro-update is executed when the corresponding guard (Fig. 4(C)) evaluates to true (e.g. `decode_itype` is executed when `instr_queue[1]`, the decoded instruction, is an I-type instruction). While `instr_queue[1]` depends on a single instruction input, in general, guards can depend on multiple inputs. For example, if `rs1_addr(instr_queue[1]) == rd_addr(instr_queue[2])`, then there is a data dependency between the current and previous instructions. In this case, the `alu_compute_alu_out_bypass_imm` micro-update is invoked (ref. Fig. 4(D)) instead of `alu_compute_rs_imm` (which uses data from the RF). While we take the example of `addi`, the model can be extended to cover other instructions. The resulting model will capture a thicker slice of the RTL with more signals, micro-updates, and guards. We illustrate example micro-updates for `sw` in Fig. 4(D).

We now discuss how this model addresses concerns from §2.

Finer temporal resolution. ISA-based approaches define semantics based on instruction-level temporal resolution, i.e., between fetch and commit points of an instruction. The micro-update model decomposes the execution of each instruction into several micro-updates, and evaluates guards and micro-updates on a per-cycle basis. Consequently, a micro-update model has a finer temporal resolution than ISA-based models by identifying the exact cycle at which a microarchitectural state update happens (blue markings in Fig. 3). This allows the micro-update model to discriminate between different microarchitectural implementations that are architecturally indistinguishable. For example, it can differentiate between the $\texttt{sw}_{\texttt{exe}}$ and $\texttt{sw}_{\texttt{mem}}$ implementations (as introduced in §2.2), and specifically identify ($\texttt{lw}; \texttt{sw}_{\texttt{exe}}$) as being vulnerable.

Imperatively modelled functional detail. While axiomatic modelling approaches (such as *μ*spec) can also provide finer temporal resolution (through microarchitectural events), they often lack functional detail. A micro-update model, however, captures functional behaviour through micro-update bodies (Fig. 4(B)). As discussed in §2.4, this allows specification of security properties like NI. Secondly, micro-updates describe imperative operations on the model state (§3.1.1). Imperative modelling enables per-cycle equivalence proofs with the RTL, which is essential for accurate model lifting. This contrasts with axiomatic approaches which require global invariants ([35, 51]).

3.1 Components of the micro-update model

3.1.1 Signals

A micro-update model captures a design slice that is identified by a set of *signals-of-interest*, which we denote as S_{data}. When generating the model, the user can specify these based on the RTL signals over which they wish to perform security analysis. At each cycle, each data signal from S_{data} is mapped to a value by an *assignment* $\sigma_{data} : \mathrm{S}_{data} \rightarrow \mathbb{V}$.

3.1.2 Micro-updates

A micro-update is an imperative code block that typically operates on small subsets of signals (Fig. 4(B)). The model as a whole comprises a set of micro-updates, L. At each cycle, signals from S_{data} are updated based on the micro-updates from L that are invoked during that cycle. The body of each micro-update $\mathsf{t} \in \mathsf{L}$ provides the semantics $[\![\mathsf{t}]\!]$, which defines how t transforms the S_{data} signals when invoked. An instruction typically performs several of these micro-updates during its lifetime, as Fig. 4(D) illustrates.

3.1.3 Trigger signals and guards

The micro-updates are invoked based on a Boolean condition called a *guard*. A guard for a micro-update t is a boolean condition evaluated over the *trigger signals* S_{trig}. We view this as a function $G_{\mathsf{t}}(\mathrm{S}_{trig}) \in \{\texttt{true}, \texttt{false}\}$. The micro-update t is executed at a cycle if and only if $G_{\mathsf{t}}(\mathrm{S}_{trig})$ evaluates to true at that cycle. Trigger signals typically correspond to top-level inputs of the RTL design (e.g. instruction inputs, memory ports). Thus the micro-update

model can directly relate micro-update invocation to the instructions inputs provided to the design. This is beneficial since it allows micro-update models to be easily connected to software-side analyses (which operate over ISA instructions).

3.1.4 Execution of the micro-update model

As introduced in §3.1.1, states of the micro-update model transition system are assignments to the signals-of-interest, $\sigma : S_{data} \rightarrow \mathbb{V}$. We denote $\Sigma_{\text{data}} = S_{data} \rightarrow \mathbb{V}$ as the set of all assignments. For initial state σ^{init} and a sequence of inputs $\vec{i}$ to the model, an execution is the sequence of states that the model is in at each cycle. We denote this as $[\![\vec{i}]\!]_{\text{m}}(\sigma^{\text{init}}) \in \Sigma^*_{\text{data}}$. Here m indicates that this is an execution (or trace) of the model. Later we will use d to indicate executions of the source RTL design.

3.2 Properties of micro-updates

Uses and modifies Knowledge of the signals used and modified by each micro-update allows us to optimize the synthesis of the micro-update model (§4.4.2). This information can be extracted from the micro-update code blocks (for example, by an AST traversal). We denote the set of signals used and modified by micro-update t as use(t) and mod(t) respectively. For example, the `lb_refill` micro-update from Fig. 4(B) has use(t) = {`mem_addr`, `mem_data`} and mod(t) = {`lb_addr`, `lb_data`, `lb_valid`}.

Sequential and combinational micro-updates We allow micro-updates of two types: combinational and sequential. These resemble sequential and combinational logic seen in Verilog-like HDLs and have similar semantics. Sequential updates consume values from the previous state, while combinational updates use values from the current cycle. Combinational updates are evaluated after dependencies have stabilized (i.e. like RTL). Presence of combinational logic in the RTL (e.g. adders) requires combinational modelling to catch the correct values.

***M*-set** While each micro-update only changes a few signals, more than one micro-update can be triggered in a single cycle of execution. We use the term *M*-set for such a set of micro-updates that can be triggered at the same cycle of execution. Not all sets of micro-updates are valid. For an *M*-set to be valid, there should not be any combinational cycles, and the constituent micro-updates must modify disjoint signals:

$$\forall \text{t}_1, \text{t}_2 \in M.\ \text{mod}(\text{t}_1) \cap \text{mod}(\text{t}_2) = \phi$$

Both these conditions can be checked from the uses and modifies information. The effect of triggering an *M*-set follows the usual sequential and combinational semantics.

4 Micro-update Model Synthesis

In this section, we discuss our framework (Fig. 5) to generate micro-update models using formal synthesis techniques. We begin by discussing the key challenges that we face.

4.1 The challenge posed by formal synthesis

While (semi)-automated approaches to synthesize formal models from RTL are useful, formal synthesis is challenging (more so than formal verification), since it requires a search over model candidates, while also verifying them for correctness. Synthesizing models involving ordering/timing constraints (e.g. for hardware) is especially challenging (compared to program synthesis [53, 67]) since the generated model must capture temporal constraints in addition to functional correctness. We navigate this complexity by decomposing the synthesis objective of a monolithic model into smaller functional (micro-updates) and temporal (guards) components.

Since individual micro-updates operate on small, local subsets of signals, identifying them is easier compared to identifying the right coordination between micro-update invocations. The latter requires a deep understanding of inter-dependencies between micro-updates. This is hard to do manually; e.g. ~10 out of 42 RTL-bug issues in the `cva6` processor [2] were due to imprecisely coordinated inter-dependencies. Our synthesis procedure does the heavy lifting of determining this choreography of micro-update invocations.

4.2 Problem formulation and guarantee

Our synthesis framework allows the user to specify the signals S_{data}, which they wish to be captured by the generated model. The choice of S_{data} exposes the tradeoff inherent to the development of formal models: models for a larger S_{data}, while more detailed, are harder to synthesize and analyze, and vice-versa. Given an RTL design, the signals S_{data}, and library L, our goal is to generate guards G_{t} for each $\text{t} \in \text{L}$, so that the behaviours of identified signals-of-interest (S_{data}) are equivalent to behaviours of corresponding signals in the source design (Property 1).

S_{data}-equivalence We now formalize the above property, which we call S_{data}-equivalence. Recall that we defined states of the micro-update model as assignments to S_{data} (§3.1.1) and executions as sequences of these assignments (§3.1.4). Similarly, we also view the source design as a transition system over all the signals from the design, S. These are a superset of signals captured in the model, $\text{S} \supseteq S_{data}$. States in the transition system of the source design are assignments to signals in S, $\sigma : \text{S} \rightarrow \mathbb{V}$. For a design state σ_{d} and a model state σ_{m}, we denote that σ_{d} and σ_{m} agree on all values in S_{data} as $\sigma_{\text{d}} \equiv_{S_{data}} \sigma_{\text{m}}$:

$$\sigma_{\text{d}} \equiv_{S_{data}} \sigma_{\text{m}} \triangleq \forall \text{s} \in S_{data}.\ \sigma_{\text{d}}(\text{s}) = \sigma_{\text{m}}(\text{s})$$

Finally, similarly to the micro-update model, we denote a trace of the design starting from state σ^{init}, on input $\vec{i}$, as $[\![\vec{i}]\!]_{\text{d}}(\sigma^{\text{init}})$. Our property states that when execution begins from initial states which the model and design agree on, and the input sequences are drawn from a set I the model continues to agree with the design for S_{data}-projected assignments at every cycle of the execution. We formalize this as follows:

Property 1 (S_{data}-Equivalence). *A model is said to satisfy S_{data}-equivalence with the source design if:*

$$\forall \vec{i} \in I.\ \forall \sigma^{\text{init}}_{\text{m}}, \sigma^{\text{init}}_{\text{d}} \in \Sigma^{\text{init}}.\ \sigma^{\text{init}}_{\text{m}} \equiv_{S_{data}} \sigma^{\text{init}}_{\text{d}} \implies$$
$$\forall j.\ [\![\vec{i}]\!]_{\text{m}}(\sigma^{\text{init}}_{\text{m}})(j) \equiv_{S_{data}} [\![\vec{i}]\!]_{\text{d}}(\sigma^{\text{init}}_{\text{d}})(j)$$

This property is guaranteed to hold because of the equivalence check performed as the final signoff on the model (§4.6).

Hyperproperty-preservation under S_{data}-equivalence The S_{data}-equivalence property allows us to infer, in a sound manner, that a program securely executes on the design if it securely executes on the model. We formalize this reasoning by recalling non-interference based security hyperproperties.

Let $S_{\text{pub}}, S_{\text{obs}}$ be subsets signals representing public inputs and (attacker)-observable outputs. A non-interference property says

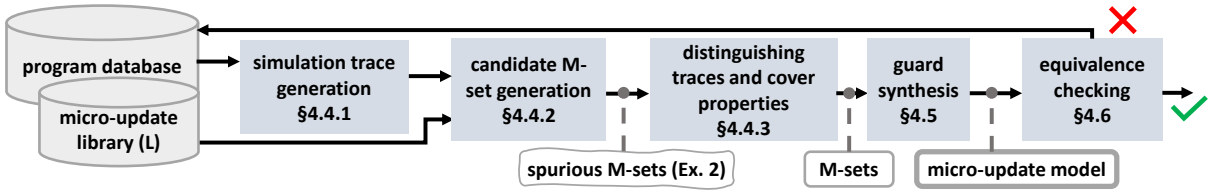

Figure 5. Synthesis pipeline for generating micro-update models.

that attacker-observable signals are only dependent on public inputs. By not being privy to non-public (secret) inputs, the attacker cannot infer sensitive information. This can be formalized as a hyperproperty [21]:

$$\mathsf{NI}_\mathsf{m}(\mathsf{S}_\mathsf{pub}, \mathsf{S}_\mathsf{obs}) \triangleq \forall \vec{i} \in I.\ \forall \sigma^1, \sigma^2 \in \Sigma^{\mathsf{init}}.$$
$$\sigma^1 \equiv_{\mathsf{S}_\mathsf{pub}} \sigma^2 \implies [\![\vec{i}]\!]_\mathsf{m}(\sigma^1) \equiv_{\mathsf{S}_\mathsf{obs}} [\![\vec{i}]\!]_\mathsf{m}(\sigma^2)$$

We refer the reader to literature (e.g., [21]) for more details. While the above property uses the model transition semantics ($[\![\cdot]\!]_\mathsf{m}$), we can also define a non-interference property NI_d by replacing the model semantics with design semantics ($[\![\cdot]\!]_\mathsf{d}$). S_{data}-equivalence implies that a non-interference property over S_{data} holds on the design if it does so on the model:

Theorem 4.1. *Suppose a model and design are S_{data}-equivalent. Then if $\mathsf{S}_\mathsf{pub}, \mathsf{S}_\mathsf{obs} \subseteq \mathsf{S}_{data}$, we have:*

$$\mathsf{NI}_\mathsf{m}(\mathsf{S}_\mathsf{pub}, \mathsf{S}_\mathsf{obs}) \implies \mathsf{NI}_\mathsf{d}(\mathsf{S}_\mathsf{pub}, \mathsf{S}_\mathsf{obs})$$

Proof sketch. We can perform a proof by contradiction. Suppose that non-interference holds on the model, but is violated on the design. Then, there exists an input sequence $\vec{i}$, and initial states $\sigma^1, \sigma^2 \in \Sigma^{\mathsf{init}}$ such that $\sigma^1 \equiv_{\mathsf{S}_\mathsf{pub}} \sigma^2$ and $[\![\vec{i}]\!]_\mathsf{d}(\sigma^1) \not\equiv_{\mathsf{S}_\mathsf{obs}} [\![\vec{i}]\!]_\mathsf{d}(\sigma^2)$. Let $\sigma^1_\mathsf{m}, \sigma^2_\mathsf{m}$ be the projections of σ^1, σ^2 on S_{data}. By S_{data}-equivalence, it follows that $[\![\vec{i}]\!]_\mathsf{m}(\sigma^1_\mathsf{m}) \equiv_{\mathsf{S}_{data}} [\![\vec{i}]\!]_\mathsf{d}(\sigma^1)$ and $[\![\vec{i}]\!]_\mathsf{m}(\sigma^2_\mathsf{m}) \equiv_{\mathsf{S}_{data}} [\![\vec{i}]\!]_\mathsf{d}(\sigma^2)$. Since $\mathsf{S}_\mathsf{pub}, \mathsf{S}_\mathsf{obs} \subseteq \mathsf{S}_{data}$, we get that $\sigma^1_\mathsf{m} \equiv_{\mathsf{S}_\mathsf{pub}} \sigma^2_\mathsf{m}$ and $[\![\vec{i}]\!]_\mathsf{m}(\sigma^1_\mathsf{m}) \not\equiv_{\mathsf{S}_\mathsf{obs}} [\![\vec{i}]\!]_\mathsf{m}(\sigma^2_\mathsf{m})$. This implies the existence of a non-interference-violating input sequence for the model, leading to a contradiction. □

We note that the above requires the signals of interest to include the observable signals. Intuitively, if one wants to explore an attack targeting a larger set of observable signals, one will need a richer model (with more signals). We demonstrate an application of this in §5.2.1 and §5.3.3, where we perform security analysis with the generated model thus inferring secure execution on the design as a result.

4.3 Synthesis overview: Figure 5

Our synthesis procedure operates in two phases. In the first phase, we generate simulation traces (§4.4.1) and use these traces to generate *candidate* M-sets (§4.4.2). These M-sets, while consistent with the simulation traces, may contain spurious candidates that do not apply generally. We use two techniques to weed out spurious candidates: distinguishing oracles and cover properties (§4.4.3). Once we eliminate these, we move to the second phase: guard synthesis (§4.5). Guard synthesis uses the M-sets from the previous phase as well as the set of trigger signals as examples when synthesizing guards for each micro-update. In this sense, guard synthesis is a form of example-based/inductive synthesis [33, 63]. Guard synthesis results in a complete candidate micro-update model. This candidate model is checked for S_{data}-equivalence (§4.6).

4.3.1 Signals-of-interest and RTL-mapping

The synthesis procedure takes as input the signals S_{data} (§3.1.1) to be included in the micro-update model, and a mapping from these signals to corresponding signals in the RTL design. The generated model then satisfies the S_{data}-equivalence property (Property 1).

4.3.2 Micro-update library

We also require the user to provide a library of micro-updates (L in Fig. 5). Candidate M-set generation (§4.4) searches over L to construct valid M-sets for the simulated traces. Since micro-updates operate on small subsets of signals, specifying this library often only requires a local understanding of the design. In our experiments, we observed that library specification time was negligible compared to design harnessing/instrumentation.

4.3.3 Discussion on the choice of micro-update library

Specificity and canonicity of micro-update library The micro-update library used when generating a model for a set of signals $S_{\mathsf{S}_{data}}$ should be able to capture all possible single-cycle transitions of the signals in $S_{\mathsf{S}_{data}}$. As such, these transitions depend on the design. For example, consider two processor designs, of which one implements bypassing while the other does not. The latter would need additional micro-updates that account for bypassing, by allowing the ALU output to be a function not only of register file outputs but also pipeline registers for stages further along in the pipeline. However, the micro-update library need not be specific to a single design. Multiple designs from the same design family can share micro-updates between them. In our experiments, we observe that the Sodor family of processors share micro-updates corresponding to instruction decode, the register file, and functional units. However, more pipelined processors require additional micro-updates to account for more intermediate stages. Finally, we note that a larger-than-necessary micro-update library does not affect the correctness of the generated model. However, as the search space increases with a larger library, it may affect the synthesis runtime. We now discuss this further.

Sensitivity of micro-update library to synthesis We now comment on the sensitivity of synthesis to the user-provided micro-update library. Firstly, libraries which have incorrect or *irrelevant* micro-updates (e.g. LB micro-updates for a model with only ALU operations) does not affect the soundness of the generated model. Such micro-updates are filtered out by candidate M-set generation, i.e., their guard is assigned `false`. Since this phase does not invoke (costly) solvers, synthesis performance is also not heavily affected.

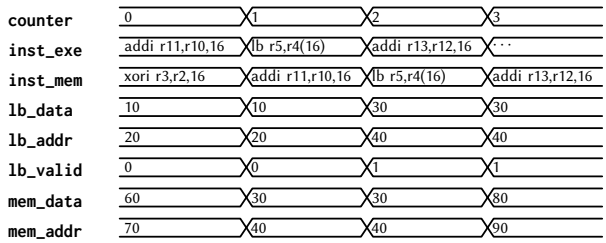

Figure 6. A sample execution trace from simulation.

Libraries which are *incomplete*, i.e. do not contain all micro-updates necessary to justify the simulation traces, fail during M-set generation. In such cases, our approach identifies the failing simulation step and signals. This helps to add missing micro-updates.

4.4 Generating M-sets

4.4.1 Simulator-based trace generation

We generate a set of random test traces of the RTL execution from a simulator (e.g. iverilog [77]). These traces must include all the signals from S_{data}. Figure 6 shows a fragment of such a trace which includes signals for the memory port and the LB component described in §2.

4.4.2 Generating candidate M-sets

A candidate M-set is a set of micro-updates that transforms the assignment to S_{data} some cycle i of a trace into the assignment at the next cycle $i+1$. As the first part of the synthesis procedure, we extract candidate M-sets that match the simulation traces. For each step of the simulation traces, we obtain a pre-post assignment pair, $(\sigma_{data}, \sigma'_{data})$. For example, from the second step in Fig. 6 we extract: $\sigma_{data} = [\texttt{lb\_data} \mapsto 10, \cdots]$ and $\sigma'_{data} = [\texttt{lb\_data} \mapsto 30, \cdots]$. The generation of candidate M-sets for $(\sigma_{data}, \sigma'_{data})$ is performed via a depth-first-search (DFS) over micro-updates from the library L. The DFS explores a sequence of micro-updates, which incrementally transforms σ_{data} into σ'_{data}. A naive DFS will consider all possible sequences of micro-updates. However, typically, several micro-updates do not depend on each other. We exploit this notion of independence to eliminate redundant search.

In general, we say that micro-update t' depends on micro-update t if t' is combinational and $\text{use}(t') \cap \text{mod}(t)$ is non-empty. Since sequential micro-updates only consume values from the previous step, they are not dependent on anything. Subsets of independent micro-updates can be searched independently in the DFS, which avoids redundant orderings.

Search pruning is also performed based on the fact that only one micro-update modifies a given signal at each step. Hence, after choosing a micro-update modifying a signal s, the DFS ignores all other micro-updates modifying s. While the number of micro-update subsets is exponential, these strategies avoid search-space explosion, making candidate M-set generation robust to the size of the micro-update library L.

Even though each candidate M-set generated by the DFS transforms the assignment σ_{data} into σ'_{data}, some of these candidates may be spurious, as we now illustrate.

Example 4.2 (Spurious M-sets)**.** Consider the trace in Fig. 6. For the third transition (`counter`=2 to 3), the `lb_refill` micro-update (Fig. 4B) correctly updates the LB signals. However, the `lb_hold` micro-update (not shown) which maintains the same values would also work for this cycle since the signals do not change. The DFS procedure will generate both of these as candidates. However, we note that in the design (Fig. 3) the LB is only refilled when the `mem` stage instruction, `inst_mem`, is a load. This is not the case for this transition since `inst_mem` is an `addi` instruction. Hence, `lb_refill` is a spurious candidate which we need to filter out.

4.4.3 Eliminating spurious M-sets

We eliminate such spurious M-sets using (a) distinguishing oracles (§4.4.3) and (b) cover properties (§4.4.3).

Using a distinguishing oracle A distinguishing oracle [39] identifies a test program (and its corresponding trace) on which two given M-sets, $M_{1,2}$, differ in their behaviour. Running M-set generation on distinguishing traces only generates one of the two M-set candidates, eliminating ambiguities such as the one in Example 4.2.

Example 4.3 (Ex. 4.2 continued)**.** If applied to the M-sets M_1 and M_2 which include micro-updates `lb_hold` and `lb_refill` respectively, a distinguishing oracle generates a trace that includes the first transition (`counter`=0 to 1) of Example 4.2. We note that only the `lb_hold` micro-update applies for this transition (since the LB entry remains invalid). The distinguishing oracle can also generate a trace with the second transition (`counter`=1 to 2), where only `lb_refill` applies. The distinguishing oracle can be used to specifically identify traces where M_1 (or M_2) is applicable and the other is not.

Adapting the technique from [39], a distinguishing oracle can be implemented by invoking a hardware model checker (e.g. SymbiYosys [20]) as follows. Given candidate M-sets M_1, M_2, we create two copies of the micro-update model. The first copy invokes micro-updates from M_1 while the second from M_2. Let Uses be signals that are used by *some* micro-update in $M_1 \cup M_2$ and Mods be signals that are modified by both M_1 and M_2. Then, we invoke a hardware model checker to generate a trace such that at some cycle i, the two copies have equal values *of all used signals*,

$$\texttt{assume}\left(\wedge_{s \in \text{Uses}}(\tau[i](s_1) = \tau[i](s_2))\right)$$

while at the next cycle $(i+1)$, M_1 and M_2 generate different values for *at least one modified signal*,

$$\texttt{assert}\left(\vee_{s' \in \text{Mod}}(\tau[i+1](s'_1) \neq \tau[i+1](s'_2))\right).$$

Since M_1 and M_2 generate different values for some signal, only one of these can be a valid candidate M-set for the transition from cycle i to $i+1$. Thus the trace generated by the model checker distinguishes M_1 and M_2.

A typical case (§5.2) where the distinguisher helps is when a functional unit operates on one of several possible inputs (e.g. in the case of bypassing). In such cases, one can identify traces in which a specific bypass path is invoked.

Using cover properties For a Boolean formula ϕ, cover properties of the form $\texttt{cover}(\phi)$ can be supplied to a hardware verification tool (e.g. JasperGold, SymbiYosys [20]). The tool, for a given cover property ϕ, aims to generate an execution in which ϕ holds. While the distinguishing oracle is useful when identifying traces on which specific M-sets do/do not apply, cover properties can be used to identify traces where certain guard predicates do/do not evaluate

to true. Traces satisfying the cover property can then be added to the trace corpus for future M-set generation. This allows the generation of rare executions that were not seen during the initial simulation.

Example 4.4 (Using cover properties)**.** Consider a simple 3-stage pipeline where the trigger signals correspond to the three instructions in the pipeline: $S_{trig} = \{i_{fet}, i_{exe}, i_{wb}\}$. Consider a bypassing path micro-update that is triggered when there is a data dependency between i_{exe} and i_{fet}, and i_{exe} does not write to zero. The guard corresponding to this micro-update is:

$$\phi_{bypass} \equiv (\texttt{rd}(i_{exe}) = \texttt{rs1}(i_{fet})) \;\&\&\; (\texttt{rd}(i_{exe}) \neq \texttt{5'b00000})$$

Since this condition requires matching register source/destination fields, the probability of it holding is somewhat low ($\sim 1/32$), and it may be missed during random simulation (§4.4.1). In such cases, the lack of example transitions for the bypassing path would result in the micro-update not being represented in the generated model. The cover property $\texttt{cover}(\phi_{bypass})$ generates an execution in which this path is triggered, mitigating this. Thus, the user can generate rare executions not observed in simulation.

4.5 Guard synthesis

The M-sets generated in the previous stage indicate micro-update invocations that specifically produce simulation traces. Guards (§3.1.3) generalize this to all traces. Intuitively, the M-sets generated in the previous stage serve as examples for guard synthesis. Guard synthesis generates guard expressions that are consistent with these examples. In this sense, guard synthesis works like programming by examples [32, 64]. Guard synthesis produces a candidate micro-update model which is then checked (§4.6) to satisfy Property 1.

We formulate guard synthesis using Syntax Guided Synthesis (SyGuS) which is a type of formal synthesis. The formal synthesis problem aims to generate an implementation of an object (typically a function) such that the generated implementation satisfies some given specification. SyGuS [6] builds on this core idea by restricting the search space of synthesis function implementations to terms from a context-free grammar and where the specification for the synthesis function is provided in the SMT-LIB format [10–12]. SyGuS is supported by multiple synthesis solvers (we use CVC5 [9]).

We formulate the guard G_{t} for micro-update t as a synthesis function over the trigger signals. For guard G_{t} we extract pairs (σ^i_{trig}, b^i) from the simulation traces and corresponding M-sets. Here, σ^i_{trig} is the trigger signal assignment at step i, and b^i is a Boolean representing whether t belongs to a valid M-set at step i. We ensure that the synthesized guard is consistent with these pairs through the following synthesis constraints.

The first constraint is that for step i, if b_i is false (meaning that micro-update t wasn't a part of any valid M-set), then the guard G_{t} must evaluate to false on σ^i_{trig}: $\bigwedge_i (b^i \vee \neg G_{\mathsf{t}}(\sigma^i_{trig}))$. If not, the generated model would be incorrect at step i.

Additionally, for each signal $\mathsf{s} \in S_{data}$, we require that exactly one micro-update modify s at each step. We define the set of micro-updates that modify s as $\mathsf{mods}(\mathsf{s})$. The following formula enforces the "exactly one" constraint for signal s:

$$\bigwedge_{\mathsf{t}\neq\mathsf{t}'\in\mathsf{mods}(\mathsf{s})} (\neg G_{\mathsf{t}}(\sigma_{trig}) \wedge \neg G_{\mathsf{t}'}(\sigma_{trig})) \wedge \bigvee_{\mathsf{t}\in\mathsf{mods}(\mathsf{s})} G_{\mathsf{t}}(\sigma_{trig}).$$

On posing this problem to a SyGuS solver, the solver returns expressions for each guard G_{t} in terms of the trigger signals. This guard expression, together with the micro-update bodies, gives us a complete model. The inability of the solver to synthesize some guard implies that the set of predicates was insufficient, i.e. there are factors outside this set that the guard depends on. The user can then try strategies such as adding more trigger signals, or increasing the SyGuS grammar depth.

4.6 Equivalence checks

Finally, the generated model is checked against the source design for S_{data}-equivalence (Property 1). If the equivalence proof fails, we get back a counterexample trace. This trace is fed back into the trace database and the synthesis loop is repeated. Future iterations of synthesis use this new trace, and hence avoid synthesizing the same (incorrect) model.

4.7 Hierarchical synthesis

The approach discussed so far monolithically generates micro-update models from the source RTL. However, our approach can benefit from the hierarchy inherent to RTL designs, which is what we now discuss. Consider a micro-update model with signals S_1 generated from a design component C_1. Now suppose we want to generate a model with signals S for a (larger) component C, of which C_1 is a sub-component. By reusing the model for C_1, we only need to generate the slice of the model corresponding to signals $S \setminus S_1$ (i.e. signals in S but not S_1). In particular, this requires generating guard predicates only for micro-updates relevant to $S \setminus S_1$. This decomposition leads to smaller synthesis queries, thus improving performance.

Hierarchical synthesis, however, requires that (a) the signals S_1 from the sub-component (C_1) not be directly driven by logic outside the component, and (b) that the trigger signals for C_1 map to trigger signals for the parent component C in a straightforward way. Here (a) ensures that S_1 signals preserve their behaviour in the larger model, and (b) ensures that the guards from C_1 in terms of trigger signals from C. We demonstrate an application of this in §5.3.2, where we first generate a model for the `store_unit` of `cva6` (C_1). Then we reuse it when generating a model for the `load_store_unit` (C), of which the the `store_unit` is a sub-component.

5 Experimental Evaluation

5.1 Methodology

We implement our framework in a Python-based tool called PAUL (Python-based Atomic-Update Language toolkit). The user specifies signals-of-interest (§4.3.1), a library of micro-updates over these signals (§4.3.2), and a set of predicates for the guards. The tool allows the user to simulate the design over input programs, generate M-sets (§4.4), and synthesize guards (§4.5).

Hyperparameters While Fig. 5 indicates the general synthesis pipeline, in practice, there are knobs that can be used to guide the synthesis. Both the distinguishing (§4.4.3) and cover (§4.4.3) trace generation require calls to a model checker, which can be time-consuming. The user can selectively apply these checks. For the distinguishing oracle, this amounts to choosing two M-sets to distinguish between, while for the cover property oracle, this amounts to providing a (possibly partial) predicate valuation.

Backend setup PAUL interfaces with iverilog [77] for simulation, and the Yosys [78] based SymbiYosys [20] for model checking. SymbiYosys, uses Boolector [56] and ABC [5] as backend solvers. We perform experimentation on an Intel Core-i7-10610U processor at 2.3GHz with 16GB RAM.

Evaluation overview We conduct two case studies by synthesizing micro-update models for: (a) `Sodor5Stage` (§5.2) and (b) components from `cva6` (§5.3.1). We demonstrate how the hierarchical lifting of micro-update models (§5.3.2) can improve scalability. We showcase applications of the lifted models for security verification of software. We demonstrate better performance compared to verification against the source RTL, and strong soundness guarantees (which are often lacking in hand-written models).

5.2 Case Study: The `Sodor5Stage` processor

The `Sodor5Stage` is a 5-stage in-order processor [71] with a simple scratchpad memory. We augment this design with the load buffer (LB) feature (as discussed in §2). We now discuss the results of model lifting, referring to Table 1.

We generate three models from the processor corresponding to different ISA subsets (rows in Tab. 1): ALUI (ALU immediate), ALUR (ALU register), and ALULS (ALU immediate + memory instructions). In each case, we identify a set of signals (S_{data}) and the library of micro-operations (L). In all models, we have a queue of the previous five instructions as the control state. Since this is a 5-stage design, this suffices to represent all inter-instruction interactions.

M-set generation In each case, we perform simulation over the respective instruction subsets by constraining the opcode. We extract the transitions and generate the candidate micro-update M-sets for these simulation traces. We observe that for the comparison instructions (e.g. `slt`), comparisons of several bypassing-path signals have the same (0/1) result leading to spurious M-sets. We provide the counts of distinguishing and cover traces used to eliminate these spurious cases in Table 1. As M-sequence generation does not involve (costly) solvers, it requires much less time (<5s) as compared to guard synthesis.

Guard synthesis and equivalence In the second phase of the synthesis, we use the M-sets to synthesize guards by formulating a SyGuS [6] query. We use cvc5 [9] as the SyGuS solver. Finally, we check the candidate micro-update model for S_{data}-equivalence with respect to the source RTL. We perform this check using a bounded model checking (BMC) query with SymbiYosys ([20]) model checker. We use a depth of 15 steps for BMC which ensures that all possible inter-instruction interactions across the pipeline are explored.

5.2.1 Do lifted models aid software security analysis?

In this section, we explore whether/how lifting micro-update models improves the trustworthiness and performance of security analysis. We begin by recalling the Bounds-Check-Bypass (BCB) gadget from the Spectre vulnerability [43, 44].

Spectre BCB background Figure 7(A) shows the Spectre BCB gadget and Fig. 7(B) shows assembly fragments for the `array2` access in the gadget. Fig. 7(C) shows the assembly fragment after a minor modification to the gadget (changing `&=` to `=`). In B (but not C), two ALU instructions separate the `lw` instruction (`101ba`) from the `sw` instruction (`101c4`). Inlays in Fig. 7 depict timing diagrams for executions of B,C on `Sodor5Stage` with the LB modification,

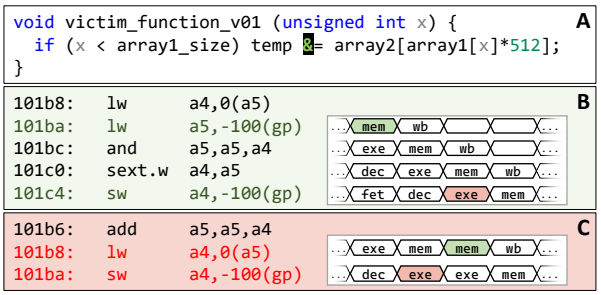

Figure 7. Subtle program changes can expose vulnerabilities: (A) Ex. 1 from Kocher's [43] examples that is vulnerable to a Spectre [44] attack. (B, C) Compiled assembly from the original example, and the example after changing the `&=` to `=`. The accompanying waveforms show how (C) remains tainted under $\texttt{sw}_{\texttt{exe}}$ while (B) does not.

under the $\texttt{sw}_{\texttt{exe}}$ implementation. In B, $\texttt{sw}_{\texttt{exe}}$ flushes the LB *after* the `lw` refills it. However, in C, $\texttt{sw}_{\texttt{exe}}$ flushes LB *before* the refill. If the `lb_addr` were *attacker observable* (as discussed in §2.2), B would be safe, while C would not.

This example shows how small changes to the microarchitecture such as the LB can render the hand-written abstract models adopted by existing approaches [18, 23, 29, 30] imprecise. As an example, the abstract model from [18] which is designed to check Spectre variants would flag both Fig. 7(B), (C) as vulnerabilities, leading to false positives. However, as we now demonstrate, the lifted micro-update model can distinguish between these outcomes and specifically flag the offending (C) case, thus resulting in more trustworthy analysis.

Semantic information-flow experiment We now demonstrate how the lifted (ALULS) model can be used to perform semantic information-flow analysis on software. We check whether there exists an information-flow from certain source (*src*) signals to the `lb_addr` signal when executing small litmus test programs. Note that `lb_addr` is assumed to be adversary-observable (§2.2). If this check passes, it guarantees that victim secrets (e.g. private keys) are not leaked through the LB side-channel when the test program is executed. We formulate information-flow as a variant of the noninterference [21] hyperproperty. We check this by invoking the SymbiYosys model checker.

Table 2 draws a comparison of the run times of these information-flow checks when performed against the lifted model and the source design. Each check (row in Tab. 2) is performed on a 1-4 instruction symbolic litmus tests (where the opcode is fixed and other fields are unconstrained). We observe that verification against the lifted model results in improved performance over the source RTL in most cases. We can guarantee the soundness of this analysis since the lifted model preserves S_{data}-equivalence with the source model (Property 1). *This demonstrates that model lifting can improve the trustworthiness of security analysis over hand-written models while providing performance improvements over the source design.*

5.3 Case Study: The `cva6` (Ariane) processor

In this section, we apply our lifting methodology to components from the `cva6` (Ariane) processor [1]. Ariane is a 6-stage application-class processor implementing the RISC-V 64-bit instruction set.

Model slice	Signals-of-interest (S_{data})	$\lvert S_{data}\rvert$	$\lvert L\rvert$	Simulation examples	Distinguish/ cover examples	Guard synthesis	Equivalence proof ($d = 15$)
I-type ALU	S_1 = decode, register file I/O, ALU, bypassing	15	25	190	1	1m1s	10m9s
R-type ALU	S_1 = decode, register file I/O, ALU, bypassing	15	29	190	3	2m14s	11m57s
ALU + `lw` + `sw`	$S_2 = S_1 \cup$ memory ports, LB	20	37	190	8	57m	34m24s

Table 1. Model parameters, example counts, synthesis and equivalence check times for the generated models for `Sodor5Stage`.

Symbolic instruction sequence (testcase)	Source signals (*src*)	Outcome Safe/Unsafe	Design runtime	Model runtime
`nop lw`	`mem`	Safe	2m43s	8s
`nop lw`	`mem, regs`	Unsafe	3m50s	11s
`nop alui lw`	`mem`	Safe	2m37s	12s
`nop sw lw`	`mem`	Safe	2m30s	9s
`nop lw lw`	`regs`	Unsafe	4m15s	8m51s
`nop( lw + alui)`^{2} `sw`	`mem, regs`	Safe	4m49s	38s
`nop( lw + alui)`^{3} `sw`	`mem, regs`	Safe	5m11s	1m31s

Table 2. Comparison of information-flow proof run times over some representative testcases. Tests are over *symbolic* sequences of instructions represented as a regular expression in the first column. The source run time is over the `Sodor5Stage` RTL (with $\texttt{sw}_{\texttt{mem}}$) while model run time is over our lifted ALULS model.

We focus on lifting four memory-facing modules in Ariane: TLB, write-buffer, store-unit, and load-store-unit. Micro-architectural units involved in memory operations (e.g. load/store buffers) are especially prone to side-channel exploits. For instance, MDS attacks (e.g. [16, 62, 72, 73]) swipe in-flight data from such buffers. Hence accurately modelling such units when performing security analysis is essential.

This case study explores two key questions: (1) in §5.3.2 we investigate whether hierarchical synthesis (§4.7) can improve the scalability of lifting and (2) in §5.3.3 we apply the lifted models to perform security analysis, and demonstrate significant improvements to verification runtime.

We now briefly describe the components we lift (for details see [1, 54]). The TLB [8, 59] is housed inside the memory management unit (MMU) of the processor and uses a PLRU (pseudo-least-recently-used) eviction scheme [58]. The write-buffer (`wbuffer`) is a coalescing buffer for the write-through data-cache, and sits between the core, data-cache and data-memory. The store unit (`store_unit`) maintains a queue of pending store requests from the core. Requests are initially queued into a speculative queue and are later moved to a commit queue (upon receiving a commit signal from the core). The load-store unit (`load_store_unit`) consists of the load and store unit sub-modules. It only handles one request at a time and stalls if there is a load-after-store conflict with an unfulfilled store operation (at the same address).

5.3.1 Generating micro-update models for `cva6`

We now discuss the generation of micro-update models for the aforementioned components, summarizing results in Tab. 3.

***TLB and* `wbuffer`** For the TLB and `wbuffer`, we generate micro-update models in which the signals-of-interest (S_{data}) include all entries in these buffers. For example, each `wbuffer` entry has three 1-bit signals representing its state: s = (`valid`, `dirty`, `txnblk`). The micro-update library L consists of five micro-updates that: (1, 2) receive a fresh write request or a replayed request to the same address, (3, 4) initiate or conclude a memory transaction, and (5) a nop (no operation). As reported in Tab. 3 while guard synthesis takes the maximum time of all the steps, the overall generation time of either model is less than ~4m.

`store_unit` *and* `load_store_unit` The load-store unit handles inter-instruction dependencies while interfacing between the core (for requests/commits) and the memory (for requests/responses). It is much more complex than the TLB and `wbuffer` components. This is reflected in the micro-update model synthesis times (Tab. 3). While the `store_unit` can be synthesized monolithically, monolithic guard synthesis for the `load_store_unit` hits a time out (✝ in Tab. 3) of 1 hr.

5.3.2 Can hierarchical synthesis improve scalability?

Since monolithic synthesis does not scale, we attempt hierarchical synthesis (§4.7) to generate a model for the `load_store_unit`. Hierarchical synthesis is feasible since the two requirements outlined in §4.7 are satisfied: (a) signals of the `store_unit` are not directly written to by outside logic and (b) almost all inputs to the (top-level) `load_store_unit` are passed to the (sub-component) `store_unit`. The exception to (b) is an input signaling whether the incoming request into the `load_store_unit` is a load/store. We conditioned it to be a store before passing it to the `store_unit`. By adopting the previously (monolithically) synthesized `store_unit`, we only needed to synthesize the non-`store_unit` guards. While monolithic synthesis timed out, we could generate the model using the hierarchical approach in ~12m (★ in Tab. 3).

Hierarchical synthesis is generally suitable when the composition is performed at module boundaries of the RTL. Such cases tend to satisfy condition (a) from §4.7. However, as seen in this case, condition (b) from §4.7 may be harder to satisfy if there is intermediate logic between the trigger signals of the parent module and those of submodules. While the logic was manually identifiable in this case, in the future one could use a version of guard synthesis to extract it. *Thus hierarchical lifting leads to better scalability when synthesizing complex models. However, this might require the user to provide additional instrumentation/hints.*

5.3.3 Is performance of security analysis better with the micro-update model compared to source RTL?

Experimental setup We investigate whether using the lifted model can improve the performance of security verification. Since we only lifted a model corresponding to the `load_store_unit`, we first wrap the generated model in a simple processor shim. This shim executes `alui`, `lw`, and `sw` instructions. While `alui` instructions are executed locally, the shim interfaces with the `load_store_unit` for `sw` and `lw` instructions. Execution is stalled if the `load_store_unit` is busy. Developing this shim required around 5-6 person-hours of work. We also wrap the source RTL

Module	Extracted signals-of-interest (S_{data})	M-set generation	Guard synthesis	Equivalence proof ($d = 15$)
TLB	states of buffer entries	<1s	4s	6s
`wbuffer`	states of buffer entries (`valid`,`dirty`,`txnblk`)	1s	2m47s	1m10s
`store_unit`	store queue, store req. states	1s	5m17s	2m58s
`load_store_unit`	store queue, store req. states, load req. state (e.g. `valid`, `spec`, `commit`, `memresp`)	2s	TO (†)	1m49s
		2s	11m38s (★)	

Table 3. Summary of the micro-update generation from the components of `cva6`. The row marked with (†) denotes monolithic synthesis, while for (★) we use hierarchical synthesis (§4.7), by reusing the `store_unit` (sub-module) generated previously.

Symbolic instruction sequence (testcase)	Constraint on the testcase program	Safe/ UnSafe	Design runtime	Model runtime
`alui sw lw alui`	NONE	US	34s	17s
`alui sw lw alui`	`regs` equal (RE)	S	3m57s	48s
`alui sw lw alui`	addr(`sw`) == addr(`lw`)	S	1m16s	34s
$\texttt{lw}_1$ `sw` $\texttt{lw}_2$ `alui`	NONE	US	1m29s	25s
$\texttt{lw}_1$ `sw` $\texttt{lw}_2$ `alui`	RE	US	1m10s	27s
$\texttt{lw}_1$ `sw` $\texttt{lw}_2$ `alui`	addr(`sw`) == addr($\texttt{lw}_2$)	S	54s	24s
$\texttt{lw}_1$ `sw` $\texttt{lw}_2$ `alui`	RE, rd($\texttt{lw}_1$) ≠ rs1(`sw`)	US	1m	28s
$\texttt{lw}_1$ `sw` $\texttt{lw}_2$ `alui`	RE, rd($\texttt{lw}_1$) ≠ rs1(`sw`) ∧ rd($\texttt{lw}_1$) ≠ rs1($\texttt{lw}_2$)	S	1m48s	44s
$\texttt{lw}_1$ `sw` $\texttt{lw}_2$ $\texttt{lw}_3$	RE, rd($\texttt{lw}_1$) ≠ rs1(`sw`) ∧ rd($\texttt{lw}_1$) ≠ rs1($\texttt{lw}_2$)	US	9m4s	1m2s
$\texttt{lw}_1$ `sw` $\texttt{lw}_2$ $\texttt{lw}_3$	RE, rd($\texttt{lw}_1$) ≠ rs1(`sw`) ∧ rd($\texttt{lw}_1$) ≠ rs1($\texttt{lw}_2$) ∧ rd($\texttt{lw}_1$) ≠ rs1($\texttt{lw}_3$)	S	12m36s	1m29s

Table 4. Security analysis checking whether the test cases result in different timing behaviours. "Design" and "Model" runtimes are for the `load_store_unit` RTL and lifted model, respectively.

for the `load_store_unit` in an identical shim for a valid comparison. We now conduct an experiment to compare security analysis run times when using the (shim-wrapped) lifted model/source `load_store_unit` RTL as the underlying HW platform (similar to §5.2.1).

The* `load_store_unit` *timing channel and results In our experiment, we analyze whether a given a software instruction sequence is vulnerable to a `load_store_unit`-based timing channel. We formulate *invulnerbility* as a non-interference [21] property stating that the timing behaviour of the instruction sequence is independent of victim data (we assume that the victim's secret resides in data memory). We check this property by invoking the SymbiYosys model checker. If the check passes, the instruction sequence is secure while a counterexample would indicate a vulnerability.

We present our results in Tab. 4, where rows denote the symbolic test cases that are checked. The test cases are based on read gadgets seen in hardware attacks. *We observe that verification run times with the lifted model are up to* 8× *lower than the source RTL.* This tends to increase for larger tests.

We also note that certain test programs are unsafe (US). This is the case since the `load_store_unit` blocks loads when there are previous pending stores at the same address. Thus, the timing behaviour of `sw` · `lw` when the `lw` and `sw` are on the same address is different than when they are on different addresses. If the address constitutes a victim secret, an attacker could infer the secret through a timing-based attack [25].

This demonstrates that models lifted from security-relevant design components can be used to check the existence of vulnerabilities in SW with significant performance improvement.

Evaluation highlights Our experiments demonstrate: (a) the feasibility of micro-update model lifting, and the ability to improve scalability through hierarchical synthesis, and (b) the application of the lifted models to perform security analysis of software with greater reliability than handwritten models and better performance than with source RTL.

6 Discussion and Limitations

Manual effort The main manual effort required in our approach is for identifying signals-of-interest (S_{data}), the micro-update library (L) and design instrumentation. The first is fundamental to our framework as S_{data} captures the design slice the user is interested in analyzing. While the current approach requires a user-specified micro-update library, going forward we foresee automating this by utilizing techniques such as specification/invariant-mining [22, 45]. As non-design-experts, a large fraction of the manual effort in our case studies, (~ 3-4 person weeks), was spent in understanding and instrumenting the RTL designs. Comparatively, the time required to write the actual micro-update library and the processor shim (in §4) was much smaller, around 1-2 person days and 4-5 person hours respectively. We believe that if lifting is performed in lockstep with the design phase, and with aid from RTL designers, the time required can be reduced.

Signals-of-interest and analysis coverage While we require the user to identify the signals-of-interest, automatically identifying the complete attack surface is a major open challenge. Prior work on analyzing hardware exploits also rely on manually identified signals (e.g. [18, 29, 30, 55, 70]). Through the signal-of-interest (S_{data}), our approach exposes a tradeoff between model coverage and scalability. While a model capturing with smaller S_{data} is less detailed, it can still be used to prove security against multiple attack scenarios targeting those signals (e.g. in §5.3.3 we study security implications of a `load_store_unit` against several instruction sequences). For generating a model with high coverage, hierarchical synthesis (§4.7) can improve scalability.

7 Related Work

Several approaches propose formal specifications/models at the ISA-level, such as the RISC-V specification [75], as well as formal models of ISA-level semantics (e.g. in SAIL [28] and Kami [41]). Approaches such as instruction-level-abstraction (ILA) [36, 37, 81, 83] extend these models to include additional architectural state (e.g. accelerator state). ILA-MCM [83] also considers a model which is

composed of code blocks, however, these are composed axiomatically. These models have been extensively used for specification [7] and verification of processor designs (e.g. [14, 61, 66, 74]). There is also work on synthesizing such models from hardware [68]. However, as discussed in §2.3, these approaches are not precise enough to model microarchitecture-level interactions, which is essential for accurate security analysis.

There is work on developing microarchitectural models of hardware, such as μspec [49]. μspec enables modelling of model memory consistency [48], coherence [52], and security properties [70] with modular and hierarchical verification [50]. More recently, there has been work on generating μspec models from RTL [35] as well as work on automatically converting axiomatic μspec models to equivalent operational models [26]. [55] uses a μspec based model to detect vulnerabilities. As discussed in §2.4, axiomatic modelling and the absence of functional detail makes checking equivalence with RTL and performing deeper specification (as required in microarchitectural security analysis) challenging.

More broadly, our work is related to *program slicing* from the software engineering domain (e.g., [34, 76]). Given a set of program variables, program slicing identifies a sub-program (*slice*) that affects the behaviours of these variables. There is also work that also performs slicing for hardware [82], where the slice captures architectural signals instead of program variables. We, however, allow the generation of models with microarchitectural signals Additionally, program slicing relies on an (often syntactic) data/control-flow analysis to identify the slice as a *subset* of the existing program. Our model, on the other hand, *replaces* the original RTL circuit with micro-updates from the user-provided library, leading to smaller and more abstract models.

Approaches that define [31] and verify [18, 23, 29, 30] security properties from a software standpoint manually develop platform models for program execution. This is error-prone due to subtle microarchitectural interactions. In §5.2.1 we demonstrated how micro-update models generated with strong guarantees can increase the level of assurance of these approaches. [80] develop a speculative platform model with a guarantee of invulnerability to some attacks, while [13] develops a methodology to validate platform models.

Unique Program Execution Checking (UPEC) [24] performs verification of RTL by checking whether executing read gadgets (e.g. [42, 44]) leads to a security vulnerability. This is orthogonal since we lift models from RTL for scalable software verification. In particular, a UPEC-like approach could be soundly performed on the lifted model owing to our equivalence guarantee (§4.2). Moreover, as our results indicate, verification against the lifted model is more performant as compared to the RTL.

Micro-updates bear resemblance to rules from BlueSpec [57] or transactions from TLM [15]. However, our focus is modelling and lifting as opposed to hardware design.

8 Conclusion

In this work, we proposed micro-update models as a formalism for developing abstract formal models of the microarchitecture. By accurately preserving security-relevant microarchitectural detail, micro-update models provide an abstract yet sound substrate for performing scalable security verification of software. To address the challenge in hand-writing formal models, we developed a semi-automated technique to hierarchically synthesize micro-update models from RTL. We evaluated our approach by synthesizing models from the `Sodor5Stage` and `cva6` processors. We demonstrated how the generated models can be used to soundly perform semantic security analysis, with improved performance over the source RTL. Our modelling and lifting framework can allow HW designers to increase the level of assurance of security analysis while reducing the efforts involved in formal model development.

Acknowledgements

We thank our reviewers and our shepherd Kristin Barber for their comments and feedback which helped improve this paper. This work benefitted from discussions with Kevin Laeufer, Federico Mora, Jonathan Shi from UC Berkeley, and Scott Constable and John Matthews from Intel Labs. This work was supported in part by Intel Corporation under the Scalable Assurance program, DARPA contract FA8750-20-C0156 and NSF grant 1837132.

References

[1] CVA6 (Ariane) User Manual. https://docs.openhwgroup.org/projects/cva6-user-manual/index.html.

[2] CVA6 RTL Bug Issues. https://github.com/openhwgroup/cva6/issues?q=is%3Aissue+label%3AType%3ABug+label%3AComponent%3ARTL+is%3Aopen. Accessed: 2023-02-20.

[3] Maaz Bin Safeer Ahmad and Alvin Cheung. Optimizing data-intensive applications automatically by leveraging parallel data processing frameworks. In *Proceedings of the 2017 ACM International Conference on Management of Data*, SIGMOD '17, pages 1675–1678, New York, NY, USA, 2017. Association for Computing Machinery.

[4] Maaz Bin Safeer Ahmad and Alvin Cheung. Automatically Leveraging MapReduce Frameworks for Data-Intensive Applications. In *Proceedings of the 2018 International Conference on Management of Data*, SIGMOD '18, pages 1205–1220, New York, NY, USA, 2018. Association for Computing Machinery.

[5] Alan Mischenko et al. Berkeley ABC tool. https://github.com/berkeley-abc/abc, 2022.

[6] Rajeev Alur, Rastislav Bodík, Garvit Juniwal, Milo M. K. Martin, Mukund Raghothaman, Sanjit A. Seshia, Rishabh Singh, Armando Solar-Lezama, Emina Torlak, and Abhishek Udupa. Syntax-guided synthesis. *2013 Formal Methods in Computer-Aided Design*, pages 1–8, 2013.

[7] Alasdair Armstrong, Thomas Bauereiss, Brian Campbell, Alastair Reid, Kathryn E. Gray, Robert M. Norton, Prashanth Mundkur, Mark Wassell, Jon French, Christopher Pulte, Shaked Flur, Ian Stark, Neel Krishnaswami, and Peter Sewell. ISA semantics for ARMv8-a, RISC-V, and CHERI-MIPS. *Proc. ACM Program. Lang.*, 3(POPL), Jan 2019.

[8] Remzi H. Arpaci-Dusseau and Andrea C. Arpaci-Dusseau. *Operating Systems: Three Easy Pieces*. CreateSpace Independent Publishing Platform, North Charleston, SC, USA, 2018.

[9] Haniel Barbosa, Clark W. Barrett, Martin Brain, Gereon Kremer, Hanna Lachnitt, Makai Mann, Abdalrhman Mohamed, Mudathir Mohamed, Aina Niemetz, Andres Nötzli, Alex Ozdemir, Mathias Preiner, Andrew Reynolds, Ying Sheng, Cesare Tinelli, and Yoni Zohar. cvc5: A versatile and industrial-strength smt solver. In *TACAS*, 2022.

[10] Clark Barrett, Pascal Fontaine, and Cesare Tinelli. The Satisfiability Modulo Theories Library (SMT-LIB). `www.SMT-LIB.org`, 2016.

[11] Clark Barrett, Roberto Sebastiani, Sanjit A. Seshia, and Cesare Tinelli. Satisfiability modulo theories. In Armin Biere, Marijn Heule, Hans van Maaren, and Toby Walsh, editors, *Handbook of Satisfiability*, chapter 33, pages 1267–1329. IOS Press, second edition, 2021.

[12] Clark W. Barrett, Roberto Sebastiani, Sanjit A. Seshia, and Cesare Tinelli. Satisfiability modulo theories. In *Handbook of Satisfiability*, 2009.

[13] Pablo Buiras, Hamed Nemati, Andreas Lindner, and Roberto Guanciale. Validation of side-channel models via observation refinement. *MICRO-54: 54th Annual IEEE/ACM International Symposium on Microarchitecture*, 2021.

[14] Jerry R. Burch and David L. Dill. Automatic verification of pipelined microprocessor control. In *Proceedings of the 6th International Conference on Computer Aided Verification*, CAV '94, pages 68–80, Berlin, Heidelberg, 1994. Springer-Verlag.

[15] Lukai Cai and Daniel Gajski. Transaction level modeling: an overview. *First IEEE/ACM/IFIP International Conference on Hardware/ Software Codesign and Systems Synthesis (IEEE Cat. No.03TH8721)*, pages 19–24, 2003.

[16] Claudio Canella, Daniel Genkin, Lukas Giner, Daniel Gruss, Moritz Lipp, Marina Minkin, Daniel Moghimi, Frank Piessens, Michael Schwarz, Berk Sunar, Jo Van Bulck, and Yuval Yarom. Fallout: Leaking Data on Meltdown-resistant CPUs. *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, 2019.

[17] Sunjay Cauligi, Craig Disselkoen, Klaus v. Gleissenthall, Dean Tullsen, Deian Stefan, Tamara Rezk, and Gilles Barthe. Constant-Time Foundations for the New Spectre Era. In *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI 2020, pages 913–926, New York, NY, USA, 2020. Association for Computing Machinery.

[18] Kevin Cheang, Cameron Rasmussen, Sanjit A. Seshia, and Pramod Subramanyan. A formal approach to secure speculation. *2019 IEEE 32nd Computer Security Foundations Symposium (CSF)*, pages 288–28815, 2019.

[19] Joonwon Choi, Muralidaran Vijayaraghavan, Benjamin Sherman, Adam Chlipala, and Arvind. Kami: A platform for high-level parametric hardware specification and its modular verification. *Proc. ACM Program. Lang.*, 1(ICFP), Aug 2017.

[20] Claire Wolf, et. al. Symbiyosys. https://github.com/YosysHQ/sby, 2022.

[21] Michael R. Clarkson and Fred B. Schneider. Hyperproperties. *2008 21st IEEE Computer Security Foundations Symposium*, pages 51–65, 2008.

[22] Michael D. Ernst, Jake Cockrell, William G. Griswold, and David Notkin. Dynamically discovering likely program invariants to support program evolution. In *Proceedings of the 21st International Conference on Software Engineering*, ICSE '99, pages 213–224, New York, NY, USA, 1999. Association for Computing Machinery.

[23] Xaver Fabian, Marco Guarnieri, and Marco Patrignani. Automatic detection of speculative execution combinations. *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, 2022.

[24] Mohammad Rahmani Fadiheh, Johannes Müller, Raik Brinkmann, Subhasish Mitra, Dominik Stoffel, and Wolfgang Kunz. A formal approach for detecting vulnerabilities to transient execution attacks in out-of-order processors. In *2020 57th ACM/IEEE Design Automation Conference (DAC)*, pages 1–6, 2020.

[25] Qian Ge, Yuval Yarom, David Cock, and Gernot Heiser. A survey of microarchitectural timing attacks and countermeasures on contemporary hardware. *Journal of Cryptographic Engineering*, 8:1–27, 2018.

[26] Adwait Godbole, Yatin A Manerkar, and Sanjit A Seshia. Automated conversion of axiomatic to operational models: Theory and practice. In *CONFERENCE ON FORMAL METHODS IN COMPUTER-AIDED DESIGN–FMCAD 2022*, page 331, 2022.

[27] Joseph A. Goguen and José Meseguer. Unwinding and inference control. *1984 IEEE Symposium on Security and Privacy*, pages 75–75, 1984.

[28] Kathryn E. Gray, Gabriel Kerneis, Dominic P. Mulligan, Christopher Pulte, Susmit Sarkar, and Peter Sewell. An integrated concurrency and core-ISA architectural envelope definition, and test oracle, for IBM POWER multiprocessors. *2015 48th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, pages 635–646, 2015.

[29] Roberto Guanciale, Musard Balliu, and Mads Dam. InSpectre: Breaking and Fixing Microarchitectural Vulnerabilities by Formal Analysis. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, CCS '20, pages 1853–1869, New York, NY, USA, 2020. Association for Computing Machinery.

[30] Marco Guarnieri, Boris Köpf, José Francisco Morales, Jan Reineke, and Andrés Sánchez. Spectector: Principled detection of speculative information flows. *2020 IEEE Symposium on Security and Privacy (SP)*, pages 1–19, 2020.

[31] Marco Guarnieri, Boris Köpf, Jan Reineke, and Pepe Vila. Hardware-software contracts for secure speculation. *2021 IEEE Symposium on Security and Privacy (SP)*, pages 1868–1883, 2021.

[32] Sumit Gulwani, William R. Harris, and Rishabh Singh. Spreadsheet data manipulation using examples. *Commun. ACM*, 55(8):97–105, Aug 2012.

[33] Sumit Gulwani, Oleksandr Polozov, and Rishabh Singh. Program synthesis. In *Software Systems Safety*, 2017.

[34] Mark Harman and Keith Brian Gallagher. Program slicing. *2008 Frontiers of Software Maintenance*, pages 58–67, 2008.

[35] Yao Hsiao, Dominic P. Mulligan, Nikos Nikoleris, Gustavo Petri, and Caroline Trippel. Synthesizing Formal Models of Hardware from RTL for Efficient Verification of Memory Model Implementations. *MICRO-54: 54th Annual IEEE/ACM International Symposium on Microarchitecture*, 2021.

[36] Bo-Yuan Huang, Hongce Zhang, Aarti Gupta, and Sharad Malik. ILAng: A Modeling and Verification Platform for SoCs Using Instruction-Level Abstractions. In Tomáš Vojnar and Lijun Zhang, editors, *Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, pages 351–357, Cham, 2019. Springer International Publishing.

[37] Bo-Yuan Huang, Hongce Zhang, Pramod Subramanyan, Yakir Vizel, Aarti Gupta, and Sharad Malik. Instruction-Level Abstraction (ILA): A Uniform Specification for System-on-Chip (SoC) Verification. *ACM Trans. Des. Autom. Electron. Syst. (TODAES)*, 24(1), dec 2018.

[38] Bo-Yuan Huang, Hongce Zhang, Pramod Subramanyan, Yakir Vizel, Aarti Gupta, and Sharad Malik. Instruction-Level Abstraction (ILA). *ACM Transactions on Design Automation of Electronic Systems (TODAES)*, 24:1–24, 2019.

[39] Susmit Jha, Sumit Gulwani, Sanjit A. Seshia, and Ashish Tiwari. Oracle-guided component-based program synthesis. *2010 ACM/IEEE 32nd International Conference on Software Engineering*, 1:215–224, 2010.

[40] Susmit Jha and Sanjit A. Seshia. A Theory of Formal Synthesis via Inductive Learning. *Acta Informatica*, 54(7):693–726, 2017.

[41] Shoaib Kamil, Alvin Cheung, Shachar Itzhaky, and Armando Solar-Lezama. Verified lifting of stencil computations. *Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation*, 2016.

[42] Vladimir Kiriansky and Carl A. Waldspurger. Speculative buffer overflows: Attacks and defenses. *ArXiv*, abs/1807.03757, 2018.

[43] Paul Kocher. Spectre Mitigations in Microsoft's C/C++ Compiler, 2018.

[44] Paul C. Kocher, Daniel Genkin, Daniel Gruss, Werner Haas, Michael Hamburg, Moritz Lipp, Stefan Mangard, Thomas Prescher, Michael Schwarz, and Yuval Yarom. Spectre attacks: Exploiting speculative execution. *2019 IEEE Symposium on Security and Privacy (SP)*, pages 1–19, 2019.

[45] Wenchao Li, Alessandro Forin, and Sanjit A. Seshia. Scalable Specification Mining for Verification and Diagnosis. In *Proceedings of the 47th Design Automation Conference*, DAC '10, pages 755–760, New York, NY, USA, 2010. Association for Computing Machinery.

[46] Moritz Lipp, Michael Schwarz, Daniel Gruss, Thomas Prescher, Werner Haas, Anders Fogh, Jann Horn, Stefan Mangard, Paul C. Kocher, Daniel Genkin, Yuval Yarom, and Michael Hamburg. Meltdown: Reading Kernel Memory from User Space. In *USENIX Security Symposium*, 2018.

[47] Fangfei Liu, Yuval Yarom, Qian Ge, Gernot Heiser, and Ruby B. Lee. Last-level cache side-channel attacks are practical. In *2015 IEEE Symposium on Security and Privacy, SP 2015, San Jose, CA, USA, May 17-21, 2015*, pages 605–622. IEEE Computer Society, 2015.

[48] Daniel Lustig, Michael Pellauer, and Margaret Martonosi. PipeCheck: Specifying and Verifying Microarchitectural Enforcement of Memory Consistency Models. *2014 47th Annual IEEE/ACM International Symposium on Microarchitecture*, pages 635–646, 2014.

[49] Daniel Lustig, Geet Sethi, Margaret Martonosi, and Abhishek Bhattacharjee. COATCheck: Verifying Memory Ordering at the Hardware-OS Interface. *Proceedings of the Twenty-First International Conference on Architectural Support for Programming Languages and Operating Systems*, 2016.

[50] Yatin A. Manerkar, Daniel Lustig, and Margaret Martonosi. RealityCheck: Bringing Modularity, Hierarchy, and Abstraction to Automated Microarchitectural Memory Consistency Verification, 2020.

[51] Yatin A. Manerkar, Daniel Lustig, Margaret Martonosi, and Michael Pellauer. RTLCheck: Verifying the Memory Consistency of RTL Designs. *2017 50th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, pages 463–476, 2017.

[52] Yatin A. Manerkar, Daniel Lustig, Michael Pellauer, and Margaret Martonosi. CCICheck: Using μhb graphs to verify the coherence-consistency interface. *2015 48th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, pages 26–37, 2015.

[53] Zohar Manna and Richard Waldinger. A deductive approach to program synthesis. *ACM Trans. Program. Lang. Syst.*, 2(1):90–121, Jan 1980.

[54] Valentin Martinoli, Yannick Teglia, Abdellah Bouagoun, and Régis Leveugle. CVA6's Data cache: Structure and Behavior, 2022.

[55] Nicholas Mosier, Hanna Lachnitt, Hamed Nemati, and Caroline Trippel. Axiomatic hardware-software contracts for security. *Proceedings of the 49th Annual International Symposium on Computer Architecture*, 2022.

[56] Aina Niemetz, Mathias Preiner, and Armin Biere. Boolector 2.0. *J. Satisf. Boolean Model. Comput.*, 9(1):53–58, 2014.

[57] Rishiyur S. Nikhil. Bluespec System Verilog: efficient, correct RTL from high level specifications. *Proceedings. Second ACM and IEEE International Conference on Formal Methods and Models for Co-Design, 2004. MEMOCODE '04.*, pages 69–70, 2004.

[58] David A. Patterson and John L. Hennessy. Computer architecture: A quantitative approach. 1969.

[59] David A. Patterson and John L. Hennessy. Computer Organization and Design, Fifth Edition: The Hardware/Software Interface. 2013.

[60] Elizabeth Polgreen, Andrew Reynolds, and Sanjit A. Seshia. Satisfiability and synthesis modulo oracles. In *VMCAI*, 2022.

[61] Alastair David Reid, Rick Chen, Anastasios Deligiannis, David Gilday, David Hoyes, Will Keen, Ashan Pathirane, Owen Shepherd, Peter Vrabel, and Ali Mustafa Zaidi. End-to-End Verification of Arm Processors with Isa-Formal. In S. Chaudhuri and A. Farzan, editors, *Proceedings of the 2016 International Conference on Computer Aided Verification (CAV'16)*, volume 9780 of *Lecture Notes in Computer Science*, pages 42–58, July 2016.

[62] Michael Schwarz, Moritz Lipp, Daniel Moghimi, Jo Van Bulck, Julian Stecklina, Thomas Prescher, and Daniel Gruss. ZombieLoad: Cross-Privilege-Boundary Data Sampling. *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, 2019.

[63] Sanjit A. Seshia. Combining induction, deduction, and structure for verification and synthesis. *Proceedings of the IEEE*, 103(11):2036–2051, 2015.

[64] Rishabh Singh and Sumit Gulwani. Learning semantic string transformations from examples. *Proc. VLDB Endow.*, 5(8):740–751, Apr 2012.

[65] Anirudh Sivaraman, Alvin Cheung, Mihai Budiu, Changhoon Kim, Mohammad Alizadeh, H. Balakrishnan, George Varghese, Nick McKeown, and Steve Licking. Packet transactions: High-level programming for line-rate switches. *Proceedings of the 2016 ACM SIGCOMM Conference*, 2016.

[66] Jens U. Skakkebæk, Robert B. Jones, and David L. Dill. Formal verification of out-of-order execution using incremental flushing. In *Proceedings of the 10th International Conference on Computer Aided Verification*, CAV '98, pages 98–109, Berlin, Heidelberg, 1998. Springer-Verlag.

[67] Armando Solar-Lezama, Liviu Tancau, Rastislav Bodik, Sanjit Seshia, and Vijay Saraswat. Combinatorial sketching for finite programs. In *Proceedings of the 12th*

International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS XII, pages 404–415, New York, NY, USA, 2006. Association for Computing Machinery.

[68] Pramod Subramanyan, Bo-Yuan Huang, Yakir Vizel, Aarti Gupta, and Sharad Malik. Template-Based Parameterized Synthesis of Uniform Instruction-Level Abstractions for SoC Verification. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 37:1692–1705, 2018.

[69] Tachio Terauchi and Alex Aiken. A capability calculus for concurrency and determinism. In Christel Baier and Holger Hermanns, editors, *CONCUR 2006 – Concurrency Theory*, pages 218–232, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.

[70] Caroline Trippel, Daniel Lustig, and Margaret Martonosi. Security Verification via Automatic Hardware-Aware Exploit Synthesis: The CheckMate Approach. *IEEE Micro*, 39:84–93, 2019.

[71] UCB-BAR. Sodor processor collection. https://github.com/ucb-bar/riscv-sodor, 2013.

[72] Stephan van Schaik, Alyssa Milburn, Sebastian Österlund, Pietro Frigo, Giorgi Maisuradze, Kaveh Razavi, Herbert Bos, and Cristiano Giuffrida. RIDL: Rogue In-Flight Data Load. *2019 IEEE Symposium on Security and Privacy (SP)*, pages 88–105, 2019.

[73] Stephan van Schaik, Marina Minkin, Andrew Kwong, Daniel Genkin, and Yuval Yarom. CacheOut: Leaking Data on Intel CPUs via Cache Evictions. *2021 IEEE Symposium on Security and Privacy (SP)*, pages 339–354, 2021.

[74] Miroslav N. Velev and Randal E. Bryant. Superscalar processor verification using efficient reductions of the logic of equality with uninterpreted functions to propositional logic. In *CHARME*, 1999.

[75] Andrew Waterman, Yunsup Lee, David A. Patterson, and Krste Asanovic. The RISC-V Instruction Set Manual. 2014.

[76] Mark Weiser. Program slicing. In *Proceedings of the 5th International Conference on Software Engineering*, ICSE '81, pages 439–449. IEEE Press, 1981.

[77] Steven Williams. iverilog. https://github.com/steveicarus/iverilog, 2022.

[78] Clifford Wolf, Johann Glaser, and Johannes Kepler. Yosys-A Free Verilog Synthesis Suite. 2013.

[79] Wenjie Xiong and Jakub Szefer. Leaking Information Through Cache LRU States. *2020 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, pages 139–152, 2020.

[80] Jiyong Yu, Mengjia Yan, Artem Khyzha, Adam Morrison, Josep Torrellas, and Christopher W. Fletcher. Speculative Taint Tracking (STT): A Comprehensive Protection for Speculatively Accessed Data. In *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*, MICRO '52, pages 954–968, New York, NY, USA, 2019. Association for Computing Machinery.

[81] Yu Zeng, Aarti Gupta, and Sharad Malik. Automatic Generation of Architecture-Level Models from RTL Designs for Processors and Accelerators. In *Proceedings of the 2022 Conference & Exhibition on Design, Automation & Test in Europe*, DATE '22, pages 460–465, Leuven, BEL, 2022. European Design and Automation Association.

[82] Yu Zeng, Bo-Yuan Huang, Hongce Zhang, Aarti Gupta, and Sharad Malik. Generating Architecture-Level Abstractions from RTL Designs for Processors and Accelerators Part I: Determining Architectural State Variables. In *2021 IEEE/ACM International Conference On Computer Aided Design (ICCAD)*, pages 1–9, 2021.

[83] Hongce Zhang, Caroline Trippel, Yatin A. Manerkar, Aarti Gupta, Margaret Martonosi, and Sharad Malik. ILA-MCM: Integrating Memory Consistency Models with Instruction-Level Abstractions for Heterogeneous System-on-Chip Verification. *2018 Formal Methods in Computer Aided Design (FMCAD)*, pages 1–10, 2018.